CS 3204
Operating Systems

Lecture 5

Godmar Back

Announcements

- Group membership due tomorrow (Sep 10), 11:59pm
- Project 1 help session tonight, 6pm-8pm, Rand 211
  - Project 1 due on Sep 29, 11:59pm
- Assigned reading on lectures page
- Please follow submission instructions
- New UTA: Nick Ryan (10h/week)

Processes & Threads

(continued)

Overview

- Have discussed:
  - User vs Kernel Mode
  - Context Switching
- Process States
- Priority Scheduling
- Process/Thread API Examples
  - Fork/join model

Process States

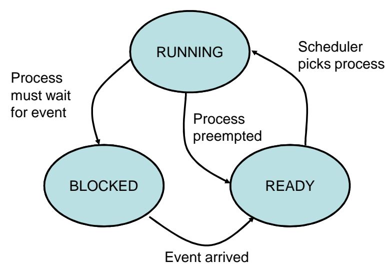

- Only 1 process (per CPU) can be in RUNNING state

Process Events

- What's an event?
  - External event:
    - disk controller completes sector transfer to memory
    - network controller signals that new packet has been received
    - clock has advanced to a predetermined time
  - Events that arise from process interaction:
    - a resource that was previously held by some process is now available (e.g., lock_release)
    - an explicit signal is sent to a process (e.g., cond_signal)
    - a process has exited or was killed
    - a new process has been created

Process Lists

- All ready processes are inserted in a “ready list” data structure
  - Running process typically not kept on ready list
  - Can implement as multiple (real) ready lists, e.g., one for each priority class
- All blocked processes are kept on lists
  - List usually associated with event that caused blocking – usually one list per object that’s causing events
- Most of scheduling involves simple and clever ways of manipulating lists

Priority Based Scheduling

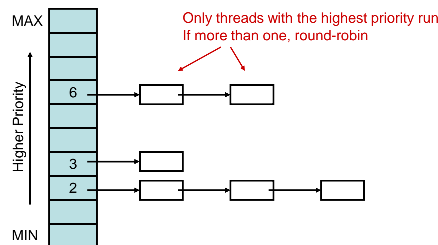

- Done in Linux (pre 2.6.23), Windows, Pintos (after you complete Project 1), …

Priority Based Scheduling (2)

- Advantage:
  - Dead simple: the highest-priority process runs
  - Q.: what is the complexity of finding which process that is?
- Disadvantage:
  - Not fair: lower-priority processes will never run
  - Hence, must adjust priorities somehow
- Many schedulers used in today’s general purpose and embedded OS work like this
  - Only difference is how/whether priorities are adjusted to provide fairness and avoid starvation
  - Exception: Linux “completely-fair scheduler” uses different scheme, will discuss that later in semester

Using thread_yield() to implement preemption

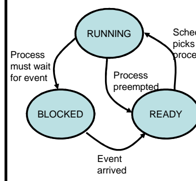

- Current thread (“RUNNING”) is moved to READY state, added to READY list.
- Then scheduler is invoked. Picks a new READY thread from READY list.
- Case a): there’s only 1 READY thread. Thread is rescheduled right away
- Case b): there are other READY thread(s)
  - b.1) another thread has higher priority – it is scheduled
  - b.2) another thread has same priority – it is scheduled provided the previously running thread was inserted in tail of ready list.
- “thread_yield()” is a call you can use whenever you identify a need to preempt current thread.
- Exception: inside an interrupt handler, use “intr_yield_on_return()” instead

Reasons for Preemption

- Generally two: quantum expired or change in priorities
- Reason #1:
  - A process of higher importance than the one that’s currently running has just become ready
- Reason #2:
  - Time Slice (or Quantum) expired
- Question: what’s good about long vs. short time slices?

I/O Bound vs CPU Bound Procs

- Processes that usually exhaust their quanta are said to be CPU bound
- Processes that frequently block for I/O are said to be I/O bound
- Q.: what are examples of each?
- What policy should a scheduler use to juggle the needs of both?

Process States w/ Suspend

(b) With Two Suspend States

- Can be useful sometimes to suspend processes
  - By user request: ^Z in Linux shell/job control
  - By OS decision: swapping out entire processes (Solaris & Windows do that, Linux doesn't)

Windows XP

- Thread state diagram in a industrial kernel

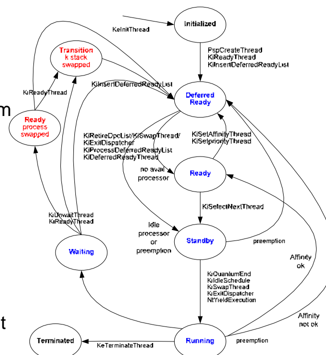

- Source: Dave Probert, Windows Internals – Copyright Microsoft 2003

Windows XP

- Priority scheduler uses 32 priorities
- Scheduling class determines range in which priority are adjusted
- Source: Microsoft® Windows® Internals, Fourth Edition: Microsoft Windows Server™

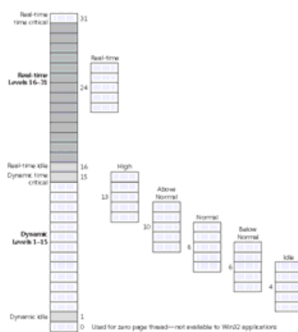

Process Creation

- Two common paradigms:
  - Cloning vs. spawning
- Cloning: (Unix)
  - "fork()" clones current process
  - child process then loads new program
- Spawning: (Windows, Pintos)
  - "exec()" spawns a new process with new program
- Difference is whether creation of new process also involves a change in program

fork()

```
#include <sys/types.h>
#include <unistd.h>
#include <stdio.h>

int main(int ac, char *av[])
{
    pid_t child = fork();
    if (child < 0)
        perror("fork"), exit(-1);
    if (child != 0) {
        printf ("I'm the parent %d, my child is  %d\n",
            getpid(), child);
        wait(NULL);    /* wait for child ("join") */
    } else {
        printf ("I'm the child  %d, my parent is %d\n",
            getpid(), getppid());

        execl("/bin/echo", "echo", "Hello, World", NULL);
    }
}
```


Fork/Exec Model

- Fork():
  - Clone most state of parent, including memory
  - Inherit some state, e.g. file descriptors
  - Important optimization: copy-on-write
    - Some state is copied lazily
  - Keeps program, changes process
- Exec():
  - Overlays current process with new executable
  - Keeps process, changes program
- Advantage: simple, clean
- Disadvantage: does not optimize common case (fork followed by exec of child)

The fork()/join() paradigm

- After fork(), parent & child execute in parallel
- Purpose:
  - Launch activity that can be done in parallel & wait for its completion
  - Or simply: launch another program and wait for its completion (shell does that)
- Pintos:
  - Kernel threads: thread_create (no thread_join)
  - exec(), you'll do wait() in Project 2

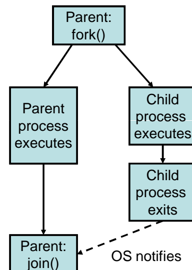

CreateProcess()

```
// Win32
BOOL CreateProcess(
  LPCTSTR lpApplicationName,
  LPTSTR lpCommandLine,
  LPSECURITY_ATTRIBUTES lpProcessAttributes,
  LPSECURITY_ATTRIBUTES lpThreadAttributes,
  BOOL bInheritHandles,
  DWORD dwCreationFlags,
  LPVOID lpEnvironment,
  LPCTSTR lpCurrentDirectory,
  LPSTARTUPINFO lpStartupInfo,
  LPPROCESS_INFORMATION lpProcessInformation );
```

- See also system(3) on Unix systems
- Pintos exec() is CreateProcess(), not like Unix's exec()

Thread Creation APIs

- How are threads embedded in a language?
- POSIX Threads Standard (in C)
  - pthread_create(), pthread_join()
  - Uses function pointer
- Java/C#
  - Thread.start(), Thread.join()
  - Java: Using "Runnable" instance
  - C#: Uses "ThreadStart" delegate
- C++
  - No standard has emerged as of yet
  - see ISO C++ Strategic Plan for Multithreading

Example pthread_create/join

```
static void * test_single(void *arg)
{
      // this function is executed by each thread, in parallel
}

/* Test the memory allocator with NTHREADS c
  pthread_t threads[NTHREADS];
  int i;
  for (i = 0; i < NTHREADS; i++)
    if (pthread_create(threads + i, (const pthread_attr_t*)NULL,
                   test_single, (void*)i) == -1)
      { printf("error creating pthread\n"); exit(-1); }

  /* Wait for threads to finish. */
  for (i = 0; i < NTHREADS; i++)
    pthread_join(threads[i], NULL);
```

Use Default Attributes – could set stack addr/size here

2nd arg could receive exit status of thread

Java Threads Example

```
public class JavaThreads {
   public static void main(String []av) throws Exception {
      Thread [] t = new Thread[5];
      for (int i = 0; i < t.length; i++) {
         final int tnum = i;
         Runnable runnable = new Runnable() {
            public void run() {
               System.out.println("Threa
            }
         };
         t[i] = new Thread(runnable);
         t[i].start();
      }
      for (int i = 0; i < t.length; i++)
         t[i].join();
      System.out.println("all done");
   }
}
```

Threads implements Runnable – could have subclassed Thread & overridden run()

Thread.join() can throw InterruptedException – can be used to interrupt thread waiting to join via Thread.interrupt

Why is taking C++ so long?

- Java didn't – and got it wrong.
  - Took years to fix
- What's the problem?
  - Compiler must know about concurrency to not reorder operations past implicit synchronization points
  - See also Pintos Reference Guide A.3.5 Memory Barriers
  - See Boehm [PLDI 2005]: Threads cannot be implemented as a library

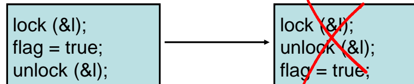

```
lock (&l);
flag = true;
unlock (&l);
```

→

```
lock (&l);
unlock (&l);
flag = true;
```